

On the Equivalence of Nonnegative Matrix Factorization and Spectral Clustering

Chris Ding[†]

Xiaofeng He[‡]

Horst D. Simon[‡]

Abstract

Current nonnegative matrix factorization (NMF) deals with $X = FG^T$ type. We provide a systematic analysis and extensions of NMF to the symmetric $W = HH^T$, and the weighted $W = HSH^T$. We show that (1) $W = HH^T$ is equivalent to Kernel K -means clustering and the Laplacian-based spectral clustering. (2) $X = FG^T$ is equivalent to simultaneous clustering of rows and columns of a bipartite graph. Algorithms are given for computing these symmetric NMFs.

1 Introduction

Standard factorization of a data matrix uses singular value decomposition (SVD) as widely used in principal component analysis (PCA). However, for many dataset such as images and text, the original data matrices are nonnegative. A factorization such as SVD contains negative entries and thus has difficulty for interpretation. Nonnegative matrix factorization (NMF) [7, 8] has many advantages over standard PCA/SVD based factorizations. In contrast to cancellations due to negative entries in matrix factors in SVD based factorizations, the nonnegativity in NMF ensures factors contain coherent parts of the original data (images).

Let $X = (x_1, \dots, x_n) \in \mathbb{R}_+^{n \times n}$ be the data matrix of nonnegative elements. In image processing, each column is a 2D gray level of the pixels. In text mining, each column is a document.

The NMF factorizes X into two nonnegative matrices,

$$X \approx FG^T, \quad (1)$$

where $F = (f_1, \dots, f_k) \in \mathbb{R}_+^{n \times k}$ and $G = (g_1, \dots, g_k) \in \mathbb{R}_+^{k \times n}$. k is a pre-specified parameter. The factorizations are obtained by the least square minimization. A number of researches on further developing NMF computational methodologies [12, 11, 10], and applications on text mining [9, 14, 11].

Here we study NMF in the direction of data clustering. The relationship between NMF and vector quantization, especially the difference, are discussed by Lee

and Seung [7] as a motivation for NMF. The clustering aspect of NMF is also studied in [14, 10].

In this paper, we provide a systematic analysis and extensions of NMF and show that NMF is equivalent to Kernel K -means clustering and Laplacian-based spectral clustering.

(1) We study the symmetric NMF of

$$W \approx HH^T \quad (2)$$

where W contains the pairwise similarities or the Kernels. We show that this is equivalent to K -means type clustering and the Laplacian based spectral clustering.

(2) We generalize this to bipartite graph clustering i.e., simultaneously clustering rows and columns of the rectangular data matrix. The result is the standard NMF.

(3) We extend NMFs to weighted NMF:

$$W \approx HSH^T. \quad (3)$$

(4) We derive the algorithms for computing these factorizations.

Overall, our study provides a comprehensive look at the nonnegative matrix factorization and spectral clustering.

2 Kernel K-means clustering and Symmetric NMF

K -means clustering is one of most widely used clustering method. Here we first briefly introduce the K -means using spectral relaxation [15, 3]. This provides the necessary background information, notations and paves the way to the nonnegative matrix factorization approach in §2.1.

K -means uses K prototypes, the centroids of clusters, to characterize the data. The objective function is to minimize the sum of squared errors,

$$J_K = \sum_{k=1}^K \sum_{i \in C_k} \|x_i - m_k\|^2 = c_2 - \sum_k \frac{1}{n_k} \sum_{i,j \in C_k} x_i^T x_j, \quad (4)$$

where $X = (x_1, \dots, x_n)$ is the data matrix, $m_k = \sum_{i \in C_k} x_i / n_k$ is the centroid of cluster C_k of n_k points,

[†]Lawrence Berkeley National Laboratory, University of California, Berkeley, CA 94720. {chqding, xhe, hdsimon}@lbl.gov

and $c_2 = \sum_i \|x_i\|^2$. The solution of the clustering is represented by K non-negative indicator vectors:

$$H = (h_1, \dots, h_n), h_k^T h_l = \delta_{kl}. \quad (5)$$

where

$$h_k = (0, \dots, 0, \overbrace{1, \dots, 1}^{n_k}, 0, \dots, 0)^T / n_k^{1/2} \quad (6)$$

Now Eq.(4) becomes $J_K = \text{Tr}(X^T X) - \text{Tr}(H^T X^T X H)$. The first term is a constant. Let $W = X^T X$. Thus $\min J_K$ becomes

$$\max_{H^T H = I, H \geq 0} J_W(H) = \text{Tr}(H^T W H). \quad (7)$$

The pairwise similarity matrix $W = X^T X$ is the standard inner-product linear Kernel matrix. It can be extended to any other kernels. This is done using a nonlinear transformation (a mapping) to the higher dimensional space

$$x_i \rightarrow \phi(x_i)$$

The clustering objective function under this mapping, with the help of Eq.(4), can be written as

$$\min J_K(\phi) = \sum_i \|\phi(x_i)\|^2 - \sum_k \frac{1}{n_k} \sum_{i,j \in C_k} \phi(x_i)^T \phi(x_j). \quad (8)$$

The first term is a constant for a given mapping function $\phi(\cdot)$ and can be ignored. Let the kernel matrix $W_{ij} = \phi(x_i)^T \phi(x_j)$. Using the cluster indicators H , the kernel K -means clustering is reduced to Eq.(7).

The objective function in Eq.(7) can be symbolically written as

$$J_W = \sum_k \frac{1}{n_k} \sum_{i,j \in C_k} w_{ij} = \text{Tr}(H^T W H). \quad (9)$$

Kernel K -means aims at maximizing within-cluster similarities. The advantage of Kernel K -means is that it can describe data distributions more complicated than Gaussian distributions.

2.1 Nonnegative factorization of Kernel K -means

We show that the optimization of Eq.(7) can be solved by the matrix factorization

$$W \approx H H^T, \quad H \geq 0. \quad (10)$$

Casting this in an optimization framework, an appropriate objective function is

$$\min_{H \geq 0} J_1 = \|W - H H^T\|^2, \quad (11)$$

where the matrix norm $\|A\|^2 = \sum_{ij} a_{ij}^2$, the Frobenius norm.

Theorem 1. $W = H H^T$ factorization is equivalent to Kernel K -means clustering with the strict orthogonality relation Eq.(5) relaxed.

Proof. The maximization of Eq.(7) can be written as

$$\begin{aligned} H &= \arg \min_{H^T H = I, H \geq 0} -2\text{Tr}(H^T W H) \\ &= \arg \min_{H^T H = I, H \geq 0} \|W\|^2 - 2\text{Tr}(H^T W H) + \|H^T H\|^2 \\ &= \arg \min_{H^T H = I, H \geq 0} \|W - H H^T\|^2. \end{aligned} \quad (12)$$

Relaxing the orthogonality $H^T H = I$ completes the proof. \square

If the nonnegativity condition is relaxed (ignored), the solution to H are the k eigenvectors with the largest eigenvalues and orthogonality is retained. Now we keep the nonnegativity of H . Will the orthogonality get lost?

Theorem 2. $W = H H^T$ factorization retains H orthogonality approximately.

proof. One can see that $\min J_1 = \|W - H H^T\|^2$ is equivalent to

$$\max_{H \geq 0} \text{Tr}(H^T W H), \quad (13)$$

$$\min_{H \geq 0} \|H^T H\|^2. \quad (14)$$

The first objective recovers the original optimization objective Eq.(7). We concentrate on 2nd term. Note

$$\|H^T H\|^2 = \sum_{lk} (H^T H)_{lk}^2 = \sum_{l \neq k} (h_l^T h_k)^2 + \sum_k (h_k^T h_k)^2.$$

Minimizing the first term is equivalent to enforcing the orthogonality among h_l : $h_l^T h_k \approx 0$. Minimizing the second term is equivalent to

$$\min \|h_1\|^4 + \dots + \|h_n\|^4. \quad (15)$$

However, H cannot be all zero, otherwise we would have $\text{Tr}(H^T W H) = 0$. More precisely, since $W \approx H H^T$,

$$\sum_{ij} w_{ij} \approx \sum_{ij} (H H^T)_{ij} = \sum_{kij} h_{ik} h_{jk} = \sum_k \|h_k\|^2, \quad (16)$$

where $\|h\| = \sum_i |h_i| = \sum_i h_i$ is the L_1 -norm of vector h . This means $\|h_i\| > 0$. Therefore, optimization of Eq.(14) with the nonzero constraint Eq.(16) implies H has near orthogonal columns, i.e.,

$$h_l^T h_k \approx \begin{cases} 0 & \text{if } l \neq k, \\ \|h_k\|^2 & \text{if } l = k. \end{cases} \quad (17)$$

Furthermore, minimization of Eq.(15) with the nonzero constraint Eq.(16) leads to the column equalization condition

$$\|h_1\| \approx \|h_2\| \approx \dots \approx \|h_n\|. \quad (18)$$

This assures the approximate balance of cluster sizes. \square

The near-orthogonality of columns of H is important for data clustering. An exact orthogonality implies that each row of H can have only one nonzero element, which implies that each data object belongs only to 1 cluster. This is hard clustering, such as in K -means. The near-orthogonality condition relaxes this a bit, i.e., each data object could belong fractionally to more than 1 cluster. This is soft clustering. A completely non-orthogonality among columns of H does not have a clear clustering interpretation.

3 Bipartite graph K -means clustering and NMF

A large number of datasets in today's applications are in the form of rectangular nonnegative matrix, such as the word-document association matrix in text mining or the DNA gene expression profiles. These types of datasets can be conveniently represented by a bipartite graph where the graph adjacency matrix B contains the association among row and column objects, which is the input data matrix $B = X$.

The above kernel K -means approach can be easily extended to bipartite graph. Let f_k be the indicator for the k -th row cluster. f_k has the same form of h_k as in Eq.(8). Put them together we have the indicator matrix $F = (f_1, \dots, f_k)$. Analogously, we define the indicator matrix $G = (g_1, \dots, g_k)$ for column-clusters.

We combine the row and column nodes together as $W = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix}$, $h_k = \frac{1}{\sqrt{2}} \begin{pmatrix} f_k \\ g_k \end{pmatrix}$, $H = \frac{1}{\sqrt{2}} \begin{pmatrix} F \\ G \end{pmatrix}$ where the factor $1/\sqrt{2}$ allows the simultaneous normalizations $h_k^T h_k = 1$, $f_k^T f_k = 1$, and $g_k^T g_k = 1$. The Kernel K -means type clustering objective function becomes

$$\max_{\substack{F^T F = I, \\ G^T G = I, \\ F, G \geq 0}} J_2 = \frac{1}{2} \text{Tr} \left(\begin{pmatrix} F \\ G \end{pmatrix}^T \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} F \\ G \end{pmatrix} \right) = \text{Tr}(F^T B G). \quad (19)$$

After some algebra, we see that J_2 minimizes the within-cluster similarities $s(R_k, C_k)$,

$$J_2 = \sum_k \frac{s(R_k, C_k)}{(|R_k| |C_k|)^{1/2}}, \quad s(R_k, C_k) = \sum_{i \in R_k} \sum_{j \in C_k} b_{ij}. \quad (20)$$

Theorem 3. The Kernel K -means clustering problem is equivalent to the following optimization problem,

$$\min_{\substack{F^T F = I, \\ G^T G = I, \\ F, G \geq 0}} J_2 = \|B - FG^T\|^2. \quad (21)$$

From Eq.(19), we have $J_2 = \text{Tr}(F^T B G)$. The optimization problem can be written

$$\begin{aligned} & \min_{\substack{F^T F = I, \\ G^T G = I, \\ F, G \geq 0}} -2\text{Tr}(F^T B G) \\ & \Rightarrow \min_{\substack{F^T F = I, \\ G^T G = I, \\ F, G \geq 0}} \|B\|^2 - 2\text{Tr}(F^T B G) + \text{Tr}(F^T F G^T G) \end{aligned}$$

In the second equation, we add two constants: $\|B\|^2$ and $\text{Tr}(F^T F G^T G) = \text{Tr}I = k$. The objective function is identical to $\|B - FG^T\|^2$. \square

Now we relax vigorous orthogonality constraints $F^T F = I; G^T G = I$ to the approximate orthogonality. Therefore, NMF is equivalent to K -means clustering with relaxed orthogonality constraints.

The orthogonality constraints play an important role. If the orthogonality holds vigorously, we can show directly that NMF of $\|B - FG^T\|^2$ is equivalent to simultaneously K -means clustering of rows and columns of B . To show this, we first prove it for clustering of columns of $B = (b_1, \dots, b_n) = (b_{ij})$, under the normalization

$$\sum_{i=1}^p b_{ij} = 1, \quad \sum_{r=1}^k g_{ir} = 1, \quad \sum_{j=1}^p f_{jk} = 1. \quad (22)$$

For any given data, normalization of X is first applied. The second normalization indicates that the i -th row of G are the posterior probabilities for b_i belonging to k clusters; they should add up to 1. The 3rd normalization $\sum_j f_{jk}$ is a standard length normalize of the vector f_k . Since we approximate $B \approx FG^T$, the normalization of FG^T should be consistent with the normalization of B . Indeed, $\sum_{i=1}^p (FG^T)_{ij} = \sum_{i=1}^p \sum_{r=1}^k F_{ir} G_{jr} = 1$, consistent with $\sum_i B_{ij} = 1$.

With this self-consistent normalization and imposing strict orthogonality on G : $g_{il}^T g_{ik} = 0, l \neq k$, we call the resulting factorization as Orthogonal NMF.

Theorem 4. Orthogonal NMF is identical to K -means clustering.

Proof. We have

$$J_2 = \|B - FG^T\|^2 = \sum_{i=1}^n \left\| b_i - \sum_{k=1}^n g_{ik} f_k \right\|^2, \quad (23)$$

because the Frobenius norm of a matrix is equivalent to sum of column norms. Now due to the normalization of G , the residue of the i -th column becomes

$$\left\| \sum_{k=1}^n g_{ik} (b_i - f_k) \right\|^2 = \sum_{k=1}^n g_{ik}^2 \|b_i - f_k\|^2 = \sum_{k=1}^n g_{ik} \|b_i - f_k\|^2$$

The 1st equality is due to the orthogonality condition of G which implies that on each row of G , only one element is nonzero. This also implies $g_{ik} = 0, 1$. Thus $g_{ik}^2 = g_{ik}$, giving the 2nd equality. Summing over i ,

$$J_2 = \sum_{k=1}^n \sum_{i \in C_k} \|b_i - f_k\|^2, \quad (24)$$

which is the standard K -means clustering of $\{b_i\}_{i=1}^n$ with f_k as the cluster centroid. ■

In NMF of optimizing $\min \|B - FG^T\|$, rows and columns are treated in equal footing, since we could equally write $J_2 = \|B^T - GF^T\|$. Thus clustering of columns of B is happening simultaneously as the clustering the rows of B .

4 Spectral clustering and NMF

In recent years spectral clustering using the Laplacian of the graph emerges as solid approach for data clustering (see references in [2]). Here we focus on the spectral clustering objective functions. There are three objectives: the Ratio Cut [6], the Normalized Cut [13], and the MinMax Cut [4]. We are interested in the multi-way clustering objective functions,

$$J = \sum_{1 \leq p < q \leq n} \frac{s(C_p, C_q)}{\rho(C_p)} + \frac{s(C_p, \bar{C}_q)}{\rho(C_q)} = \sum_{k=1}^n \frac{s(C_k, \bar{C}_k)}{\rho(C_k)} \quad (25)$$

$$\rho(C_k) = \begin{cases} |C_k| & \text{for Ratio Cut} \\ \sum_{i \in C_k} d_i & \text{for Normalized Cut} \\ s(C_k, \bar{C}_k) & \text{for MinMax Cut} \end{cases} \quad (26)$$

where \bar{C}_k is the complement of subset C_k in graph G , $s(A, B) = \sum_{i \in A} \sum_{j \in B} w_{ij}$, and $d_i = \sum_j w_{ij}$.

Here we show that the minimization of these objective functions can be equivalently carried out via the non-negative matrix factorizations. The proof follows the multi-way spectral relaxation[3] of Normalized Cut and MinMax Cut. We focus on Normalized Cut.

Theorem 5. Normalized Cut using pairwise similarity matrix W is equivalent to Kernel K-means clustering with the kernel matrix

$$\tilde{W} = D^{-1/2} W D^{-1/2}, \quad (27)$$

where $D = \text{diag}(d_1, \dots, d_n)$.

Corollary 5. Normalized Cut using similarity W is equivalent to nonnegative matrix factorization

$$\min_{H \geq 0} J_3 = \|\tilde{W} - HH^T\|^2. \quad (28)$$

Proof of Theorem 5. Let h_k be the cluster indicators as in Eq.(6). One can easily see that

$$s(C_k, \bar{C}_k) = \sum_{i \in C_k} \sum_{j \in \bar{C}_k} w_{ij} = h_k^T (D - W) h_k \quad (29)$$

and $\sum_{i \in C_k} d_i = h_k^T D h_k$. Define the scaled cluster indicator vector $z_t = D^{1/2} h_t / \|D^{1/2} h_t\|$, which obey the orthonormal condition $z_t^T z_k = \delta_{tk}$, or $Z^T Z = I$, where $Z = (z_1, \dots, z_n)$. Substituting into the Normalized Cut objective function, we have

$$J_{NC} = \sum_{t=1}^n \frac{h_t^T (D - W) h_t}{h_t^T D h_t} = \sum_{t=1}^n z_t^T (I - \tilde{W}) z_t$$

The first term is a constant. Thus the minimization problem becomes

$$\max_{Z^T Z = I, Z \geq 0} \text{Tr}(Z^T \tilde{W} Z) \quad (30)$$

This is identical to the Kernel K-means clustering of Eq.(7). Once the solution \tilde{Z} is obtained, we can recover H by optimizing

$$\min_{H \geq 0} \sum_t \left\| z_t - \frac{D^{1/2} h_t}{\|D^{1/2} h_t\|} \right\|^2. \quad (31)$$

The exact solution are $h_k = D^{-1/2} z_k$, or $H = D^{-1/2} Z$. Thus row i of Z is multiplied by a constant $d_i^{-1/2}$. The relative weight across different cluster in the same row remain same. Thus H represents the same clustering as Z does. □

Theorem 5 show the spectral clustering are directly related to Kernel K-means clustering, which is equivalence to NMF by Theorem 1. Thus NMF, Kernel K-means clustering and spectral clustering are unified in a simple way: they are different prescriptions of the same problem with slightly different constraints.

5 Weighted Nonnegative $W = HSH^T$

In both Kernel K-means and spectral clustering, we assume the pairwise similarity matrix W are semi positive definite. For kernel matrices, this is true. But a large number of similarity matrices is nonnegative, but not s.p.d. This motivates us to propose the following more general NMF:

$$\min_H J_5 = \|W - HSH^T\|^2, \quad (32)$$

When the similarity matrix W is indefinite, W has negative eigenvalues. HH^T will not provide a good approximation, because HH^T can not absorb the subspace associated with negative eigenvalues. However, HSH^T can absorb subspaces associated with both positive and negative eigenvalues, i.e., the indefiniteness of W is passed on to S . This distinction is well-known in linear algebra where matrix factorizations have Cholesky factorization $A = LL^T$ if matrix A is s.p.d. Otherwise,

one does $A = LDL^T$ factorization, where the diagonal matrix D takes care of the negative eigenvalues.

The second reason for nonnegative $W = HSH^T$ is that the extra degrees of freedom provided by S allow H to be more closer to the form of cluster indicators. This benefits occur for both s.p.d. W and indefinite W .

The third reason for nonnegative $W = HSH^T$ is that S provides a good characterization of the quality of the clustering. Generally speaking, given a fixed W and number of clusters K , the residue of the matrix approximation $J_0^{\text{opt}} = \min \|W - HSH^T\|^2$ will be smaller than $J_1^{\text{opt}} = \min \|W - HH^T\|^2$. Furthermore, the K -by- K matrix S has a special meaning. To see this, let us assume H are vigorous cluster indicators, i.e., $H^T H = I$. Setting the derivative $\partial J_0 / \partial S = 0$, we obtain

$$S = H^T W H, \text{ or } S_{lk} = h_l^T W h_k = \frac{\sum_{j \in C_l} \sum_{j \in C_k} w_{lj}}{\sqrt{n_{lk}}} \quad (33)$$

S represents properly normalized within-cluster sum of weights ($l = k$) and between-cluster sum of weights ($l \neq k$). For this reason, we call this type of NMF as weighted NMF. The usefulness of weighted NMF is that if the clusters are well-separated, we would see the off-diagonal elements of S are much smaller than the diagonal elements of S .

The fourth reason is the consistency between standard $W = HH^T$ and $B = FG^T$. Since we can define a kernel as $W = B^T B$. Thus the factorization $W \approx B^T B \approx (FG^T)^T (FG^T) = G(F^T F)G^T$. Let $S = F^T F$, we obtain the weighted NMF.

6 Symmetric NMF Algorithms

We briefly outline the algorithms for computing symmetric factorizations $W = HH^T$ and $W = HSH^T$. For $W = HH^T$, the updating rule is

$$H_{ik} \leftarrow H_{ik} \left(1 - \beta + \beta \frac{(WH)_{ik}}{(HH^T H)_{ik}} \right). \quad (34)$$

where $0 < \beta \leq 1$. In practice, we find $\beta = 1/2$ is a good choice. A faster algorithm¹

$$H \leftarrow \max(WH(H^T H)^{-1}, 0). \quad (35)$$

can be used in the first stage of the iteration. Algorithmic issues of symmetric NMF is also studied in [1].

For weighted NMF $W = HSH^T$, the update rules are

$$S_{ik} \leftarrow S_{ik} \frac{(H^T WH)_{ik}}{(H^T H SH^T H)_{ik}}. \quad (36)$$

¹For the nonsymmetric NMF of Eq.(1), the algorithm is $F \leftarrow \max(BG(G^T G)^{-1}, 0)$, $G \leftarrow \max(B^T F(F^T F)^{-1}, 0)$. Without nonnegative constraints, these algorithms converge respectively to global optimal solutions of J_1 in Eq.(11) and J_2 in Eq.(21).

$$H_{ik} \leftarrow H_{ik} \left(1 - \beta + \beta \frac{(WHS)_{ik}}{(HSH^T HS)_{ik}} \right). \quad (37)$$

Acknowledgement. This work is supported by U.S. Department of Energy, Office of Science, under contract DE-A03-76SF00098 through a LBNL LDRD grant.

References

- [1] M. Catral, Lixing Han, Michael Neumann, and Robert J. Plemmons. On reduced rank nonnegative matrix factorizations for symmetric matrices. *Linear Algebra and Its Applications*, to appear.
- [2] C. Ding. A tutorial on spectral clustering. *Int'l Conf. Machine Learning (ICML 2004)*, 2004.
- [3] C. Ding and X. He. K-means clustering and principal component analysis. *LBNL-52988. Int'l Conf. Machine Learning (ICML 2004)*, 2004.
- [4] C. Ding, X. He, H. Zha, M. Gu, and H. Simon. A min-max cut algorithm for graph partitioning and data clustering. *Proc. IEEE Int'l Conf. Data Mining*, 2001.
- [5] M. Gu, H. Zha, C. Ding, X. He, and H. Simon. Spectral relaxation models and structure analysis for k-way graph clustering and bi-clustering. *Penn State Univ Tech Report CSE-01-007*, 2001.
- [6] L. Hagen and A.B. Kahng. New spectral methods for ratio cut partitioning and clustering. *IEEE Trans. on Computer Aided Design*, 11:1074–1085, 1992.
- [7] Daniel D. Lee and H. Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999.
- [8] Daniel D. Lee and H. Sebastian Seung. Algorithms for non-negative matrix factorization. In T. G. Dietterich and V. Tresp, editors, *Advances in Neural Information Processing Systems*, volume 13. The MIT Press, 2001.
- [9] Stan Z. Li, Xinwen Hou, Hongjiang Zhang, and Qiansheng Cheng. Learning spatially localized, parts-based representation. In *Proceeding of 2001 IEEE Computer Vision and Pattern Recognition*, pages 207–212, Kauai, Hawaii, 2001.
- [10] Tao Li and Sheng Ma. IFD: Iterative feature and data clustering. In *Proceedings of the 2004 SIAM International conference on Data Mining (SDM 2004)*, pages 472–476, Lake Buena Vista, Florida, 2004.
- [11] V. Paul Pauca, Farid Shahnaz, Micheal W. Berry, and Robert J. Plemmons. Text mining using non-negative matrix factorization. In *Proceedings of the 2004 SIAM International conference on Data Mining (SDM 2004)*, pages 452–456, Lake Buena Vista, Florida, 2004.
- [12] Fei Sha, Lawrence K. Saul, and Daniel D. Lee. Multiplicative updates for non-negative quadratic programming in support vector machines. In S. Thrun S. Becker and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 1041–1048. MIT Press, Cambridge, MA, 2003.
- [13] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 22:888–905, 2000.
- [14] Wei Xu, Xin Liu, and Yihong Gong. Document clustering based on non-negative matrix factorization. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 267–273, Toronto, Canada, 2003.
- [15] H. Zha, C. Ding, M. Gu, X. He, and H.D. Simon. Spectral relaxation for K-means clustering. *Advances in Neural Information Processing Systems 14 (NIPS'01)*, pages 1057–1064, 2002.